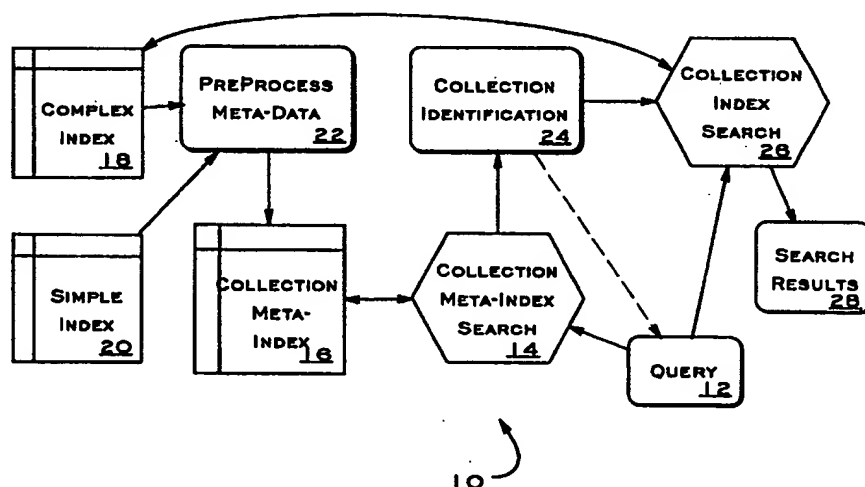


INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁶ : G06F 17/30	A1	(11) International Publication Number: WO 99/14691 (43) International Publication Date: 25 March 1999 (25.03.99)									
<p>(21) International Application Number: PCT/US98/18844</p> <p>(22) International Filing Date: 10 September 1998 (10.09.98)</p> <p>(30) Priority Data:</p> <table border="0"> <tr> <td>08/928,542</td> <td>12 September 1997 (12.09.97)</td> <td>US</td> </tr> <tr> <td>08/928,543</td> <td>12 September 1997 (12.09.97)</td> <td>US</td> </tr> <tr> <td>08/928,294</td> <td>12 September 1997 (12.09.97)</td> <td>US</td> </tr> </table> <p>(71) Applicant: INFOSEEK CORPORATION [US/US]; 1399 Moffett Park Drive, Sunnyvale, CA 95089-1134 (US).</p> <p>(72) Inventors: KIRSCH, Steven, T.; 916 Lundy Lane, Los Altos, CA 94024 (US). CHANG, William, I.; 1631 Mercy Street, Mountain View, CA 94041 (US).</p> <p>(74) Agents: LOVEJOY, David, E. et al.; Fliesler, Dubb, Meyer and Lovejoy LLP, Suite 400, Four Embarcadero Center, San Francisco, CA 94111-4156 (US).</p>		08/928,542	12 September 1997 (12.09.97)	US	08/928,543	12 September 1997 (12.09.97)	US	08/928,294	12 September 1997 (12.09.97)	US	<p>(81) Designated States: AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, CA, CH, CN, CU, CZ, DE, DK, EE, ES, FI, GB, GE, GH, GM, HU, ID, IL, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MD, MG, MK, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, UA, UG, UZ, VN, YU, ZW, ARIPO patent (GH, GM, KE, LS, MW, SD, SZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG).</p> <p>Published With international search report.</p>
08/928,542	12 September 1997 (12.09.97)	US									
08/928,543	12 September 1997 (12.09.97)	US									
08/928,294	12 September 1997 (12.09.97)	US									

(54) Title: METHODS FOR ITERATIVELY AND INTERACTIVELY PERFORMING COLLECTION SELECTION IN FULL TEXT SEARCHES



(57) Abstract

A method of selecting a subset of a plurality of document collections for searching in response to a predetermined query is based on accessing a meta-information data file that describes the query significant search terms that are present in a particular document collection correlated to normalized document usage frequencies of such terms within the documents of each document collection. By access to the meta-information data file, a relevance score for each of the document collections is determined. The method then returns an identification of the subset of the plurality of document collections having the highest relevance scores for use in evaluating the predetermined query. The meta-information data file may be constructed to include document normalized term frequencies and other contextual information that can be evaluated in the application of a query against a particular document collection.

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece	ML	Mali	TR	Turkey
BG	Bulgaria	HU	Hungary	MN	Mongolia	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MR	Mauritania	UA	Ukraine
BR	Brazil	IL	Israel	MW	Malawi	UG	Uganda
BY	Belarus	IS	Iceland	MX	Mexico	US	United States of America
CA	Canada	IT	Italy	NE	Niger	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NL	Netherlands	VN	Viet Nam
CG	Congo	KE	Kenya	NO	Norway	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NZ	New Zealand	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	PL	Poland		
CM	Cameroun	KR	Republic of Korea	PT	Portugal		
CN	China	KZ	Kazakstan	RO	Romania		
CU	Cuba	LC	Saint Lucia	RU	Russian Federation		
CZ	Czech Republic	LI	Liechtenstein	SD	Sudan		
DE	Germany	LK	Sri Lanka	SE	Sweden		
DK	Denmark	LR	Liberia	SG	Singapore		
EE	Estonia						

TITLE:

5 Methods for Iteratively and Interactively Performing
Collection Selection in Full Text Searches

INVENTORS:

10 Steven T. Kirsch
William I. Chang

BACKGROUND OF THE INVENTIONField of the Invention:

15 The present invention is generally related to full text document
searching and retrieval, as may be performed over local and wide-area networks, and
in particular to a method of performing effective document searches over multiple,
independent document collections

Description of the Related Art:

20 During the past few years, the quantity and diversity of information and
services available over the public (Internet-type) and private (Intranet-type) local and
wide area networks, generically referred to as the "Internet," has grown substantially.
In particular, the variety of information accessible through such Internet based services
is growing rapidly both in terms of scope and depth.

25 One of the significant benefits of information being accessible over the Internet
is that very diverse information can be accessed in a largely presentation independent
form. A number of independent Internet search services exist to provide context
based, content derived indexes searchable over the Internet through a query based

interface. Consequently, much if not all of the diverse information available through the Internet can be found and utilized by individuals and companies, or simply "users," who use the Internet.

5 While access to much of the information available through the Internet is free for public use, numerous proprietary or fee-based access document collections exist. Although such private document collections may all be accessible through the Internet, which is increasingly preferred over the many existing proprietary modem networks, the collections are generally forced to be accessible as discrete entities in order to maintain fee-based access control. In effect, such private document collections are
10 restricted to use on a collection access for fee basis.

Private document collections are likely to continue to exist as significant sources of unique information. Independent content creators and providers derive significant revenues from the licensing of private collection content typically to collection access providers who, in turn, derive revenue from fee-based access by
15 users to various available combinations of private collections.

In order to maximize the desirability for users to access a particular private collection and preferably related sets of private collections, a collection access provider will acquire licensed rights to make available a wide variety of individual collections of content related documents as discrete databases that can be manually
20 selected for search by a user. Typically, searches and retrievals of information from the discrete databases are subject to specific access fees determined based on the relative commercial worth of the information maintained in the individual databases. Consequently, access fees are typically calculated on the number of documents that are variously searched, reviewed, and retrieved in preparation of a search report from
25 a particular database.

A known problem in providing access to multiple databases is the relative difficulty or inefficiency in identifying an optimal database or set of databases that should be searched to obtain the best search report for a some particular unstructured,

or *ad hoc*, database query. In order to support even the possibility of *ad hoc* queries, the database search must be conducted on a full text or content established basis. Existing full text search engines typically allow a user to search many databases simultaneously. For example, commercial private collection access providers, such as Dialog, allow a user to search some 500 or more different databases either individually or in manually selected sets. Consequently, the selection of a most appropriate set of databases to search places a substantial burden on the user for each query. The user must manually determine and select a particular set of databases that must, by definition, contain the desired results to a query. Such a database set selection is difficult since the selection is made preemptively and independent of the query. This burden may be even more of an issue where access fees are charged for conducting a search against a database even where no search responsive documents are found or examined. In the aggregate, this problem is typically referred to as the "collection selection problem."

The collection selection problem is complicated further when the opportunity and desire exists to search any combination of public and private document collections. The Internet effectively provides the opportunity to access many quite disparately located and maintained databases. The importance of solving the selection collection problem thus derives from the user's desire to ensure that, for a given *ad hoc* query, the best and most comprehensive set of possible documents will be returned for examination and potential use at minimum cost.

The collection selection problem is formidable even when dealing with a single collection provider. Dialog, an exemplary collection access provider, alone provides access to over 500 separate databases, many with indistinct summary statements of scope and overlapping coverage of topics. With over 50,000 databases estimated presently available on the Internet, the collection selection problem is therefore impractical to solve reliably and efficiently by a user.

Some approaches to providing automated or at least semi-automated solutions

to the collection selection problem have been developed. Known techniques, such as WAIS (wide area information server), utilize a "server of servers" approach. A "master" database is created to contain documents that describe the contents of other "client" databases as may be potentially available on the Internet. A user first selects
5 and searches the master database to identify a set of client databases that can then be searched for the best results for a given query.

In many instances, a master WAIS database is constructed and updated manually. The master database can also be generated at least semi-automatically through the use of automatons that collect information freely from the Internet. The
10 performance of such automatons, however, is often imperfect, if not simply incorrect, in their assessments of client databases. Even at best, certain client databases, including typically private and proprietary document collections, may block access by the automatons and are thus completely unrepresented in the master database.

Even where database access can be obtained and document summaries
15 automatically generated, the scaling of the master database becomes problematic if only due to the incomplete, summary, and mis-characterization of document summary entries in the master database. Manual intervention to prepare and improve automaton generated document summaries will enhance the usefulness of the master database. When any manual intervention is required, however, the scaling of the
20 master database comes at least at the expense of the useful content of the master database document summary entries. With greatly increased scale, often only abbreviated document titles or small fractions of the client database documents can be collected as summaries into the master database. As scale increases, succinct manually generated summaries of client database documents become increasingly
25 desired, if not required, to provide any adequate content for the master database document entries. Unfortunately, even at only a modest scale, a master database of manually generated or modified document summaries becomes an impracticable construct to build or maintain.

Perhaps one of the most advanced scalable approaches to constructing and using a meaningful master database is a system known as *GLOSS* (Glossary-of-Servers Server). An automaton is typically used to prepare a master database document for each client database that is to be included within *GLOSS*. Each master database document effectively stores the frequency of whatever potential query terms occur within the corresponding client collection of documents. The master database documents are then stored as the master records that collectively form the master database.

In response to a user query, *GLOSS* operates against the master database documents to estimate the number of relevant client collection documents that exist in the respective client collections. These relevant document estimates are determined from a calculation based on the combined query term frequencies within each of the master database documents. *GLOSS* then assumes that client databases ranked as having the greatest number of combined query term occurrences are the most relevant databases to then search.

Unfortunately, utilizing a relevance system based on term frequency inherently constrains the type and effectiveness of queries that can be meaningfully directed against the master database. In addition, the estimator used by *GLOSS* is by definition a specific to any client document. The *GLOSS* system is therefore highly subject to failures to identify client databases that may contain only a relatively few instances of the query terms, yet may contain relevant documents.

Other approaches to establishing a quantitative basis for selecting client database sets includes the use of comprehensive indexing strategies, ranking systems based on training queries, expert systems using rule-based deduction methodologies, and inference networks. These approaches are used to examine knowledge base descriptions of client document collections.

Indexing and ranking systems both operate typically against the client databases directly to, in effect, create categorizations of the client databases against

search term occurrences. All possible query terms are indexed in the case of comprehensive indexing, while a limited set of predefined or static query terms are used in the case of simple ranking. Indexing thus generates a master database of selectable completeness that is nonetheless useable for selecting a most likely relevant set of client databases for a particular query. Ranking also generates a master database, though based on the results of a limited set of broad test queries intended to collectively categorize subsets of the available client databases. In effect, categorization by fixed query term results in generally orthogonal lists of ranked client database sets.

Expert system approaches typically operate on client database scope and content descriptions to deduce or establish a basis for subsequently deducing a most likely set of databases that will likely contain the most relevant documents for a particular query.

Finally, inference networks utilize a term-frequency based probabilistic approach to estimating the relevance of a particular client database as against other client databases. Unfortunately, the known implementations of inference networks are unable to accurately rank the potential relevance of client databases of diverse size and differences in the generation of summaries for each of the client databases considered.

Thus, the known approaches to solving the client database collection selection problem are generally viewed as inefficient in the assembly, construction, and maintenance of a master document database. These known systems are also viewed as often ineffective in identifying the likely most relevant documents within entire sets of collections because real world collections are often highly variable in size, scope, and content or cannot be uniformly characterized by existing quantitative approaches.

Another and perhaps practically most significant limitation of these known systems is that each must be self-contained in order to operate. This is a direct result of each system utilizing a proprietary algorithm, whether implemented as a manual operation or through the operation of an automaton, to universally assemble the

information necessary to create or populate the master database documents from the raw collection documents. As such, these known systems cannot depend on one-another or on any other indexing systems; each must be responsible for both the total generation and subsequent exclusive utilization of their master database summary record documents.

Consequently, there is a clear need for an enhanced system of handling the collection selection problem in view of the ever increasing number and scale of collections available on the Internet and the increasing variety of the collections, both in terms of existing organization and informational content.

SUMMARY OF THE INVENTION

Thus, a general purpose of the present invention is to provide an efficient and effective method of selecting the likely most relevant collections for searching based on an *ad hoc* query.

This is achieved in the present invention by providing a method of selecting a subset of a plurality of document collections for searching, in response to a predetermined query, based on accessing a meta-information data file that correlates the query significant search terms present in a particular document collection with normalized document usage frequencies of such terms within the documents of each document collection and a normalized document collection frequency of documents that include the search significant terms within the set of document collections. By access to the meta-information data file, a relevance score for each of the document collections is determined. The method then returns an identification of the subset of the plurality of document collections having the highest relevance scores for use in evaluating the predetermined query.

An advantage of the present invention is that the method provides for both automated and manual description to be used in selecting collections that contain the

most likely relevant documents in relation to an *ad hoc* query.

Another advantage is that both hard boolean selection terms, such as "*document type is newspaper*," may be combined with soft terms, such as "*choose the 10 collections containing the most articles about 'Bill Gates'*."

5 A further advantage of the present invention is that the method makes use of term proximity, capitalization, phraseology and other information in establishing a relevance ranking of the collections with respect to an *ad hoc* query.

Still another advantage of the present invention is that the method can operate to create and maintain a meta-index database based on whatever information may be
10 gleaned from existing proprietary database indexes for the collections that may be searched.

Yet another advantage of the present invention is that the method may operate from a meta-index database created from standardized format collection statistics summary files autonomously developed from the individual collections and preferably
15 from the collection indexes. By using the collection indexes directly as the source of information for the summary files, the present invention establishes an independence from the proprietary algorithms used in creating the source collection indexes.

A still further advantage of the present invention is that the construction of the collection statistics summary files from collection indexes and, in turn, the
20 construction of the meta-index database from the statistics summary files can and preferably does preserve substantial collection and document specific information such as, but not limited to collection name, document type, authors, date of publication, publisher, keywords, summary description of contents, price, language, country of publication, publication name, number of documents in the collection, the total size
25 of the collection, the average document size and average number of words in the base document collection for use in evaluating collection relevance to a potentially complex and qualified *ad hoc* query.

Still another advantage of the present invention is that a number of different

ranking algorithms can be used to identify the likely most relevant collections depending in-part potentially on the particular summary information available for each collection considered. The ranking algorithms used may also variably use proximity, capitalization and other contextual information if and as preserved through the construction of the meta-index database, thereby maximizing the efficiency of the present invention in identifying the likely most relevant collections to search.

A yet further advantage of the present invention is that utilization of the collection indexes, with their basic document location and term content data, as the source information for the summary files allows per document and collection statistics to be gathered. These statistics allow documents that have a significant likelihood of relevance to be identified in collections of documents that are otherwise seemingly irrelevant to a particular *ad hoc* query.

BRIEF DESCRIPTION OF THE DRAWINGS

These and other advantages and features of the present invention will become better understood upon consideration of the following detailed description of the invention when considered in connection of the accompanying drawings, in which like reference numerals designate like parts throughout the figures thereof, and wherein:

Figure 1 is a flow diagram illustrating the operation of a preferred embodiment of the present invention in supporting a meta-index database construction and user search;

Figure 2 is flow diagram illustrating the generation of a multi-tiered meta-index composed of multiple data types from either or both of existing collection text and prepared collection indexes in accord with a preferred embodiment of the present invention;

Figure 3 is a flow diagram illustrating the tiered query operation of a preferred embodiment of the present invention in efficiently and cost effectively producing a ranked set of document collections that are most likely to be relevant to the object of

the query;

Figure 4 is a flow diagram illustrating the operation of a preferred embodiment of the present invention in generating and updating a meta-index composed of multiple data types;

5 Figure 5 is a detailed diagram illustrating the preferred and alternate distributed system operations that support automated and semi-automated collection selection and client collection searching; and

 Figure 6 provides a simplified illustration of a multi-function collection search engine constructed in accordance with a preferred embodiment of the present
10 invention.

DETAILED DESCRIPTION OF THE INVENTION

An Internet-based computerized document collection search methodology 10 is generally illustrated in Figure 1. The methodology 10 is preferably implemented,
15 in a preferred embodiment of the present invention, as an application program or cooperative set of application programs executing on one or more conventional networked computer systems. In operation, a user query 12 is used to select a qualified set of most likely relevant document collections. This collection selection operation is performed by a meta-index collection selection search engine 14 against
20 a master index database 16, preferably referred to as a collection meta-index. The user query 12 is then used by a collection index search engine 26 against the identified collection set to select some number of most likely relevant documents from the identified collection set. A listing of the identified collection documents then forms the search report 28 that is then returned in response to the user query 12.

25 In greater detail, a user query 12 may be preferentially formed using any combination of a free text syntax, such as "growing orchids indoors," "new Sun products," and "a computer system with a RISC processor in it," fielded text qualifiers, and various pre- and post-search qualifiers that can be used to limit the

scope of the collections to be searched and selected. Any qualifiers and descriptors available to the meta-index collection selection search engine 14 can be used. Typical pre-search qualifiers can define document language (English), document type (research journal; not press releases), and cost (search free; document retrieval less than \$0.10). The resulting set of qualified document collections, as determined by the collection meta-index search engine 14, is then used in evaluating the free text portion of the query 12. Application of post-search qualifiers, such as the maximum number of collections to select for searching, can be deferred. Explicit user identification of qualifiers as pre- or post-search is unnecessary, since such identification can be inferred from the type of the qualifier.

The free text user query 12 is then selectively processed utilizing a space character as a default search term separator. The query processor also operates to remove or conform internal punctuation, such that "I.B.M." and "IBM" as well as "\$1,000" and "\$1000" are equivalently recognized. Common variably represented terms, such as online, are also equivalently recognized; online being made equivalent to "on line" and "on-line" for purposes of term equivalence mapping. Capitalization and other indirect hints, as well as direct or explicit hints, such as a "+" or "-" term prefix to require or negate existence of a term in the query results, may be used to variously augment and refine the meaningful context of the user query. The various search terms are thus reduced and expanded as appropriate into a list of specifically identified search terms.

For purposes of collection selection resolution, these identified terms are not then stemmed, preserving potential prefix and suffix contextual information. Preferably, search queries are parsed so that "PDP-11/44M+," "Internet-in-a-box," and "SVR4.0.2.1," as well as "options" (preserving context in connection with "stock options") and "greater" (preserving context in connection with "greater Bay Area"), are all identified as individual and distinct search terms for purposes of collection set selection.

Punctuation specified within a definable list, such as "[{<=>},.:!/?'", is used to define punctuation and potentially other characters that are preferably ignored as term delimiters where the punctuation or other listed character occurs anywhere as part of a search term.

5 In accordance with the present invention, some identified search terms, though otherwise proper, are also selectively ignored. These terms, largely conjunctions and other context-free words, are generally referred to as stop-terms. A stop-list of stop-terms is defined for use by the query processor in processing the query 12. By recognizing and discarding context free stop-terms, the individual search terms of context laden phrases, such as "computer system RISC processor" from "a computer
10 system with a RISC processor in it," can be identified and manipulated to maintain proximity and relative order information at least for purposes of initially selecting a collection set for searching. As indicated, at least relative proximity information is maintained.

15 The fielded text portion of a query 12 may include categorically predefined qualifier terms as explicit search limiters. Various applicable predefined fielded search qualifier terms, such as "Collection Type," "Author," "Publication Name," "Publication Date," and others can be inferentially recognized by the query processor. Where a qualifier term is recognized, a corresponding fielded search term limitation is
20 combined with the free text portion of the query through the use of defined boolean and relative relation operators.

The effective search terms of the processed query 12 are then provided to the collection meta-index search engine 14 for analysis against a collection meta-index 16. The meta-index search engine 14 may be any conventional index based search engine.
25 The collection meta-index 16 itself is preferably formed as a compilation of documents where each document stores meta-index data derived from the index or indexes that describe a corresponding base collection. The meta-index 16 thus contains a set of documents that directly correspond to the set of document collections potentially

searchable in response to any user query 12.

The collection meta-index 16 documents can be prepared through a preprocessing 22 of base collection indexes 18, 20, often referred to generically as meta-data, that were previously and independently generated by various different indexing algorithms to describe corresponding base document collections. The particular algorithms used in generation of the indexes 18, 20, which may have fundamentally different structures and informational type content, are not of particularly essential importance to the use and operation of the present invention. Preferably, the meta-data of the indexes 18, 20 are directly preprocessed 22 to produce meta-index documents, also referred to as collection summary records, of standardized format. Information characteristically (language), if not explicitly (cost), describing the collection is stored in the respective summary records as fielded text or data.

Thus, the preferred standardized summary record structure preserves a combination of fielded data, term frequencies for contextually distinctive search terms, and proximity information relating the various search terms indexed. A collection summary record may be generated by either a collection content provider or a collection access provider, though the collection content provider will have more immediate access to the base collection indexes, knowledge of the specific structure of the base collection's index files, and knowledge of the specific documents added to the base collection since any prior generation of a corresponding summary record structure.

Preferably, the summary record structure is or will be standardized for use by all collection access providers who may provide access to particular base collections. By utilizing standardized summary record structures, the base collection content providers have a standardized basis for supporting collection searching independent of the search algorithms utilized by any particular content access provider. Similarly, the standardized structure allows content access providers to accept collection

summary information from any number of different collection content providers independent of any particular indexing algorithms utilized in creating specific base collection indexes.

5 Fields within the summary record structure provide for the storage of fundamental collection statistics, fielded data lists, a term list, and term proximity relations. The fundamental collection statistics include, directly or indirectly, the definitive name of the collection and information regarding the collection content provider, an optional summary description of the collection, the number of documents within the collection, the number of terms, preferably excluding stop-terms, that occur
10 in the collection, the number of occurrences of each term in each documents of the collection, the number of times each term appears in the collection, and the length of each document in the collection. The fielded data lists include a data entry for each of the documents in the collection for each type of fielded data. The term list provides a comprehensive, ordered listing of all of the terms, excluding stop-terms, that occur
15 in the collection correlated to the documents of their occurrence. Finally, the term proximity relations identify term phrases that occur in corresponding documents of the collections. The term phrases are consecutively ordered search terms, excluding stop-terms and that do not span a sentence terminator, fixed in sets of two or more terms as they occur in the documents of a base collection. In a preferred embodiment of the
20 present invention, term phrases can be chosen to be short series of two or possibly three terms.

Where collection summary records are prepared by the collection content providers, or perhaps by a third party service company who operates on behalf of some group of collection content providers, each collection summary record can be
25 pushed, preferably using a secure Internet protocol, to each of the existing authorized collection access providers. The summary records can be prepared and pushed to the collection access providers on at least an as needed basis to reflect significant updates in the contents of a base collection. Each time a collection access provider properly

receives an updated summary record, their collection meta-index is correspondingly updated and any prior existing summary record is overwritten or deleted.

Alternately, the content access providers may pull new and updated summary records from base collection content providers. Again, the actual transfer of the summary records is preferably by a secure Internet protocol. This allows the collection content providers to potentially derive revenue based on the rate of update pulls and therefore the currency of the summary record information that any particular content access provider receives.

The content access provider may directly utilize the collection summary records to create collection summary records for the meta-index 16. However, in a preferred embodiment of the present invention, the base collection summary records are further processed by specific collection access providers separately or in parallel with the generation of the base collection indexes to optimize the organization of the collection meta-index to any particular search algorithms implemented by the access provider as part of the collection meta-index search engine 14. Specifically, multiple sets of meta-index documents may be generated for any given collection. For example, meta-index documents can be constructed to reflect the base documents themselves, to reflect any internally referenced documents, and multiple subject matter covered by a document. By installing all of the meta-index documents produced from the summary records in the collection meta-index 16, each meta-index document provides a separate basis for considering the relative relevance of a set of base collection documents.

The collection meta-index search engine 14 operates on the collection meta-index 16 to identify a particular set of base collections that contain documents that are likely most relevant to the user in response to the query 12. That is, the collection meta-index search engine identifies and ranks through a ranking algorithm an ordered list of individual meta-index documents, representing respective collections, that are considered to be the most relevant to the search terms of the user query 12. Since

each meta-index document identifies a corresponding base document collection, a ranked set of collections is then determined as likely containing the most relevant documents to the user's query 12.

5 In operating on the collection meta-index 16, the collection meta-index search engine 14 may further limit the meta-index documents or preselect against those documents under consideration to those that meet the fielded data requirements of the user query 12 as qualified by whatever applicable boolean and logical operators are specified in the query 12. Since the fielded data is preferably typed and organized in a uniform structure, the collection meta-index search engine 14 may perform a simple
10 linear search against the meta-index documents to initially locate appropriately conforming base collections.

Alternately, the collection meta-index 16 may itself be indexed at least on the fielded text and data. In this case, the collection meta-index search engine may perform an index based query to identify those documents represented in the
15 collection meta-index that again meet the presently applicable fielded text and data requirements of the user query 12.

The remaining free text derived portion of the query 12 and query specific fielded text qualifiers are then used to evaluate the likely relevance of any remaining collection meta-index documents to the user query 12. The search terms are selected
20 against the fielded text, term, and phrase lists of the meta-index documents. This selection results in a ranking of the meta-index documents in order of the likely most relevant to the likely least relevant in view of the specific user query 12.

The list of likely most relevant collections is then identified 24 from the list of likely most relevant meta-index documents that are within a predetermined or specific
25 user defined maximum number of collections to select at one time for searching, or that meet at least a minimum relevance ranking threshold, or both.

In a preferred embodiment of the present invention, the list of selected collections is then effectively combined with the user query 12 and submitted to a

collection index search engine 26. Alternately, subject to user preference definition, a new or refined query 12' can be constructed by the user to redefine the collection selection already made or to provide a new or refined query 12" to the collection index search engine 26 for use against the presently selected document collection set as
5 determined by the initial query 12 or a query 12'. This supports a tiered approach to collection set identification and searching of the identified collection set. That is, the list of selected collections may be presented back to the user for manual adjustment and approval prior to being submitted with the initial query 12 or a subsequent query 12" to the collection index search engine 26. Consequently, the present invention
10 supports automatic, semi-automatic, and tiered collection selection resolution along with flat and hierarchical queries without requiring the user to select base collections to search beyond selection of a particular collection access provider.

The collection index search engine 26 may be a conventional search engine configured to operate against the collection indexes specified by the list of selected
15 base collections. Search results 28 are then produced in accordance with the conventional operation of the collection index search engine 26. In accordance with an alternate embodiment of the present invention, the collection meta-index search engine 14 and collection index search engine 26 may be one and the same search engine. In this instance, the single search engine 14, 26 is selectable, perhaps by the
20 specification of the meta-index or indexes to search, to perform either a meta-index search in accordance with the present invention or a conventional base collection index search.

The present invention thus significantly provides for the creation and subsequent update of a summary record comprehensively for each base collection, not
25 on a per collection access provider basis. Thus, each collection content provider need only prepare and update a single collection summary record for distribution to all authorized collection access providers.

In addition, the provision of a standardized summary record to the collection

access providers still allows each collection access provider to apply their own proprietary algorithms for computation of applicable search term weights, term stemming, and proximity weights for each term of a collection. A collection access provider is thus not bound by fundamental or even extraneous aspects of the proprietary algorithms utilized by a collection content provider in generating any particular base collection index. Since the standardized summary records are constructed to carry a substantial amount of fundamental information about a collection, generally beyond that presently used by conventional search engines, including unstemmed terms, hyphenated terms, capitalization and proximity information, the collection access providers are not hindered in any way by limitations on the form and content of the data provided by the proprietary general base summary records. Indeed, the preferred summary records of the present invention are advantageously generated during or in parallel with the base collection indexing process, rather than being extracted solely from existing base collection inverted indexes and other summary information.

Finally, the present invention allows combining the results of multiple meta-collection searches to create a final likely most relevant collection ranking where multiple intersecting user queries are needed or desired in producing a final combined search result.

Figure 2 provides an overview of the process used to generate a collection meta-index. Source collection data, in the form of collection text and prepared standard and proprietary collection indexes, as available, is processed to define as necessary the collection summary fielded text and data, the statistical data, and term lists. This information is then combined through a potentially proprietary process step to finally produce the collection meta-index.

The preferred collection selection process is shown in Figure 3. The initial user query is received and constrained by any pre-search constraints applicable to the universe of collections currently available through the user selected

Internet access provider. A constrained set of collections 50 is thus formed. This set of collections 50 may be no more than a logical identification of the collection summary records stored by the collection meta-index 16 representing the universe of collections 46.

5 From the constrained collections 50, a document frequency term is calculated 52 based on the free text and un-applied fielded text provided as part of the user query 12. Document frequency (d_F), for purposes of the present invention, is defined as a normalized percentage of documents in the current collection (d_{C1}) that have at least one occurrence of the term (t). A minimum threshold (Th) of documents in the current
10 collection is enforced to preclude collections containing less than Th documents from dominating the eventual ranking of the collections. A small number, such as 100, is chosen for the value of Th . The document frequency is preferably calculated as given in Eq. 1. The preferred normalization is performed using logarithmic scaling.

$$d_F = \log \left[\frac{d_{C1,t}}{(d_{C1} > Th: d_{C1} \geq Th)} \cdot (.01) \right] \quad \text{Eq. 1}$$

15 Next, an inverse collection frequency term is then calculated 54 based on the constrained collection set 50 and the free text and same un-applied fielded text used in calculation of the document frequency (d_F). The inverse collection frequency (ICF) is defined in Eq. 2 as a normalized value given by the total number of documents in the collection set divided by the total number of documents in the collection set that have at least one occurrence of the term (t). To maintain proportionality with the calculated
20 document frequency (d_F), the inverse collection frequency is also preferably normalized using a logarithmic scale. However, the calculation is performed over the full set of the constrained collections ($d_{C\alpha}$). The ICF therefor needs to be calculated once for a particular constrained collection 50.

$$ICF = \log \left[\frac{\sum_{x=1}^n d_{Cx}}{\sum_{x=1}^n d_{Cx;t}} \right] \quad \text{Eq. 2}$$

A rank is then determined for each collection within the set of constrained collections 50. The collection rank R for a current collection (d_{C1}) and term (t) is defined in Eq. 3.

$$R_{(C1,t)} = d_{F(C1)} \cdot ICF_{(t)} \quad \text{Eq. 3}$$

- 5 Where collection ranking involves multiple terms (t), a final rank R_F for a particular collection d_{Cx} is the sum of the rank values $R_{(Cx,t)}$ for each term (t). Thus, calculation of a final rank value R_F for a particular collection is defined in Eq. 4.

$$R_{F;Cx} = \sum_{t=1}^{tn} R_{Cx;t} \quad \text{Eq. 4}$$

- 10 The post-search constraints, such as maximum number of collections returned and minimum level of relevance, are then applied to the rank ordered set of collection rankings. This produces an at least logical identification of the ranked base collections 60 determined to be most likely relevant to responding to the original user query 12.

- 15 Figure 4 provides a more detailed view of the preferred steps 62 used in preprocessing meta-data or, alternately, the underlying base collection document data, to create a base collection meta-index 16 document describing the contents of a corresponding base document collection. To begin, a determination is made 64 as to whether some pre-existing meta-data description of a collection is suitable for use in creating a sufficient summary record. Where no other access to the collection documents is available, the collection meta-data is presumed sufficient, even where substantial information is absent. Where access to the collection is not limited to just

the meta-data, the meta-data may be determined to be sufficiently complete in preserving descriptive attributes of the collection, including collection descriptive statistics, capitalization, distinctive terms, distinctive phrases or sufficient location information to reconstruct occurrences of distinctive phrases, and proximity information relating to the location of terms as they appear in the collection documents, so that an acceptable meta-index summary record can be derived from the base collection meta-data.

Where the derivation of the summary record from the meta-data requires some substantial reconstructive work, the summary record may be better and easier derived from the underlying base collection documents either in combination with the generation of the collection meta-data or as a separate process. This is particularly true where the collection content provider undertakes the generation of the base collection summary record.

Once the source of sufficient descriptive information is identified 64, a basis for establishing term densities for document specific terms is determined generally from the lengths of the documents within the collection. Preferably, a normalized document length is determined 66 either by choice by the collection content provider based on knowledge of the collection or by empirical default as the document size exceeded by approximately 80% of the documents within the collection. For a collection composed substantially of published articles and related documents, the normalized document size may be around ten typical publication paragraphs. Where the collection is substantially composed of technical specifications, such as the ANSI specifications, the normalized document size will be far larger.

The term densities of specific terms can then be calculated against the collection's normalized document length 68. A generic stop-list is used to exclude potential terms that do not carry contextual weight. Preferably, the generic stop-list is augmented for a particular collection with terms that may be deemed too common to provide significant basis for distinguishing content between the documents of that

particular base collection. For example, in a collection of pharmaceutical reports, terms such as "FDA," "drugs," and other similar terms that are extensively used in that collection might be added to the stop-list. Conversely, "in-vitro" would not be added to the stop-list unless substantially all documents within the collection are known, typically by pre-analysis of the base collection itself, to concern "in-vitro" technology. In most instances, the collection content provider is in the best position to identify terms that frequently appear in the collection documents, yet do not serve to accurately characterize or distinguish the contents of the base documents within the collection from one another.

Term densities are then determined for the remaining terms that occur within the collection. Of course, other algorithms or strategies for further identifying terms of likely significance within the collection can also be used to limit the number of terms where term densities are determined. By basing term density determinations on normalized document lengths, the present invention provides a significant sensitivity to documents that address a certain aspect of a subject even where the remainder of the documents in the collection are silent on this aspect. Terms that are used frequently even in only a single document of the collection are thus still identified by a high term density. A list of significant specific terms is then created with a term order or ranking based on relative term density 70.

In addition to the specific terms identified by substantial relative term densities, general terms, consisting of single words and multi-word phrases, are considered 71 for their potential to characterize documents within the collection. Again subject to the use of a stop-list, the densities of identifiable general terms within the base collection are determined. Additional qualifications on the use of the stop-list include not excluding instances of stop-list words that are hard-hyphenated with non-stop list words or that are capitalized, but do not appear at the beginning of a sentence. Preferably, phrases are limited to contiguously occurring sets of two terms, i.e., not including any intervening stop-list words, though phrases of longer lengths could be

considered. Also preferably, two term phrases that are frequently part of three or more term phrases are not eliminated as general terms. This allows relatively contextually important words to be identified through their overlapping occurrence as single word terms and as part of multiple word phrases. Thus, a natural increased weighing is derived for the term "computer" where the term is used alone and as part of "computer program" and "RISC computer." The general terms are then formed into a list with an order or ranking based on relative general term densities 71.

A summary record for the collection as a whole can then be created 70, 71, 72, 74, 76 by combination of the various information that is recorded in a summary record. The specific term list is evaluated 70 to identify those specific terms that are most significant. Significance may be determined by the establishment of a minimum or threshold term density. Significance may also be determined by selecting some number of the specific terms that have the highest relative term densities.

The general term list is evaluated 71 to also select out the most significant general terms for the collection. Again, significance may be determined on the basis of general term density threshold or as some number of general terms that have the highest relative term densities. In a preferred embodiment of the present invention, the significant specific and general terms are determined interdependently. The number of significant terms is limited to a combined total of a set number of terms, e.g. terms appearing in at least 1% of the documents or possibly a fixed number such as 5,000. These terms are selected from a combined list of both specific and general terms ordered by term densities. These significant terms, existing as a collection content dependant combination of single word terms and term phrases, are then used as part of the summary record 70, 71.

In addition to storing the significant terms, various other data is stored directly with or in direct correspondence to the significant terms. This additional data may include relative term proximity information and express term weights calculated based on each term's frequency of occurrence within the base collection as a whole.

Finally, various information that serves to categorically 72, statistically 74, and otherwise 76 describe the collection is added to the summary record. The category information includes all of the fielded data that is specific to the individual documents within the collection or of the collection as a whole. In a preferred embodiment of the present invention, the fielded data stored for individual documents and the collection as a whole includes, but is not limited to, collection name, document type, document title, authors, date of publication, publisher, keywords, summary description of collection or document, document length, price, language, country of publication, and publication name. In addition, the preferred stored collection specific statistical data includes, but is not limited to, the number of documents in the collection, the total size of the collection, the average document size and average number of words in the base document collection.

The summary record information 70, 71, 72, 74, 76 is then formatted into a concise standardized record format. The resulting standardized summary record 78 can then be added directly as a document to a collection meta-index 16. The summary record can also be further processed at the discretion of each particular content access provider to a proprietary format and then added as a document to their own collection meta-index 16.

In Figure 5, a complete multiple collection system 80 that supports automated collection selection resolution and searching over multiple independent collections is shown. A user, represented by a computer workstation 82, generates a query 84 that utilizes the Internet 86 as a transport mechanism. The query 84 conventionally is directed against any one of the three illustrated collections 88, 90, 92 where the collections are presented through independent collection access providers. Where one or more of the collections 88, 90, 92 are accessible through a common content access provider, the user query 84 is preconditioned upon issuance from the user workstation 82 to specify the particular set of collections 88, 90, 92 that the query is to be resolved against. Also conventionally, index processors 94, 96, 98 are utilized to create the

respectively proprietary indexes 100, 102, 104 that are used in execution of the query 84 execution against the collections 88, 90, 92.

In accordance with the present invention, the individual collection indexes 100, 102, 104 are transported over the Internet 86 or by other secure means to a meta index processor 106. The indexes 100, 102, 104 may be the ordinary indexes utilized separately in connection with accesses of the collections 88, 90, 92. Alternately, the collection content providers or others who control the index processors 94, 96, 98 may choose to generate separate indexes for the collections 88, 90, 92 that directly conform to the desired content definition and structure desired by the meta index processor 106. Where the existing indexes 100, 102, 104 are to be utilized, the collection content provider or meta index searcher, who operates the meta index processor 106, conforms to the greatest extent reasonable the content and structure of the indexes 100, 102, 104 to that desired by the meta index processor.

The meta index processor 106 operates on the provided indexes to generate multiple meta indexes 108, 110, 112, 114 that are stored in a meta index data base 118. The individual collection indexes 116, as restructured for the meta index processor 106 are also stored by the meta index database 118.

The various meta indexes 108, 110, 112, 114, each provide, in effect, a view to the contents of the combined collections 88, 90, 92. Thus, the query 84, presented without necessary precondition as to the particular set of collections to search is passed through the Internet 86 to a search engine 120 operated by the meta-index searcher. As shown in Figure 6, the search engine 120 preferably includes a preference selectable and definable interface 122 for receiving a query 84 and managing the collection selection process. Where the query 84 allows or requires collection selection resolution, a meta index search engine 124 is provided with the query 84 for use and operation against the meta indexes 108, 110, 112, 114 stored by the meta index database. The meta index search engine 124 produces a ranked listing of the most relevant collections 88, 90, 92 that could be utilized in satisfying the query 84. The

collections are ranked based on a relevancy score calculated base on the number of search term matches.

Depending on the preferences established in connection with servicing the query 84, the search interface 122 may refine the collection selection list produced by the meta index search engine 124 based on a number of preference criteria. These criteria may include limits on the number of collections to be concurrently searched, a minimum relevance score, collection access cost constraints and the like. The search interface 122 may also provide for a transaction over the Internet 86 to the user workstation 82 to allow the end user to variously review, refine and confirm the collection selection list. The opportunity to confirm the collection selection list is itself preferably just another preference that can be established within the search interface 122 from the user workstation 82.

Once a confirmed collection selection list is obtained, the query 84 is preferably executed by an index search engine 124 against the indexes 116 selectively limited to those indexes corresponding to the collection selection list. The index search engine 120 thus obtains a scored and ranked list of documents from the collections 88, 90, 92 within the collection selection list that are responsive to the query 84. The search interface 122 may again refine this collection document list based on preferences established from the user workstation 82. These preferences may define refinement criteria such as minimal acceptable relevancy score, total number of documents identified, maximum number of documents from a particular collection, and the like. The refined document list is then provided by the search interface 122 through the Internet 86 as a report 121 to the user workstation 82. Thus, the end user can analyze the report 121 and selectively retrieve identified documents from the various collections 88, 90, 92 as desired.

Thus, a method for providing for the automated or selectively semi-automated selection of document collections to be searched in response to a particular user query has been described. While the present invention has been described particularly with

reference to Internet accessible document collections, the present invention is equally applicable to the searching of non-documentary collections that can be treated as document collections accessible through a network other than the Internet and specifically including proprietary and/or private intranets.

5 In view of the above description of the preferred embodiments of the present invention, many modifications and variations of the disclosed embodiments will be readily appreciated by those of skill in the art. It is therefore to be understood that, within the scope of the appended claims, the invention may be practiced otherwise than as specifically described above.

10

15

20

CLAIMS

- 1 1. A method permitting iterative performance of collection selection
2 relative to a set of databases, where each said database includes a plurality of
3 documents, to obtain consistent relative-ranking collection selection results for each
4 iteration, said method comprising the steps of:
- 5 a) obtaining a collection selection query including a set of
6 predetermined search terms;
- 7 b) determining an inverse collection frequency for each member of said
8 set of predetermined search terms with respect to each said database and said set of
9 databases, and determining a document frequency for each member of said set of
10 predetermined search terms with respect to each said database;
- 11 c) determining a ranking value for each said database based on a sum
12 of the products of said inverse collection frequencies for said set of predetermined
13 search terms and said document frequencies for respective members of said set of
14 search terms;
- 15 d) selecting a subset of said set of databases based on predetermined
16 criteria dependant on said ranking value for each said database; and
- 17 e) selectively repeating portions of said steps (b) through (d) with
18 respect to each member of said set of predetermined search terms for each iteration of
19 said method.
- 1 2. The method of Claim 1 wherein an iteration of said method is
2 performed upon modification of said set of predetermined search terms.
- 1 3. The method of Claim 2 wherein modification of said set of
2 predetermined search terms includes, in any combination, the addition, alteration, and
3 removal of a search term with respect to said set of predetermined search terms.

1 4. The method of Claim 2 further comprising a step of preparing a meta-
2 index including a set of records respectively representative of said set of databases,
3 each member of said set of records including any of search term data, statistical data
4 and fielded data, said inverse collection frequencies and said document frequencies
5 being determinable from said search term data and said statistical data.

1 5. The method of Claim 4 wherein said set of databases is selected from
2 a superset of databases based on a set of conditional values established in connection
3 with said set of predetermined search terms, the selection of said set of databases being
4 based on an intersection of said set of conditional values and said fielded data stored
5 with respect to said superset of databases.

1 6. The method of Claim 5 wherein an iteration of said method is
2 performed upon modification of said set of conditional values.

1 7. The method of Claim 6 wherein said search term data includes single
2 word terms and multiple-word phrase terms and wherein said set of predetermined
3 search terms includes in any combination of one or more:
4 a) a search term including a single-word term;
5 b) a search term including one or more single word terms reduced by
6 exclusion of context-free single word terms;
7 c) a search term identified as a noun phrase;
8 d) a search term identified as a word tuple;
9 e) a search term identified as a capitalized single word term or a
10 multiple word term;
11 f) a proper name reduced by exclusion of any context-free single word
12 terms;

- 13 g) a hyphenated phrase expanded to include, as separate search terms,
14 the single-word terms of the hyphenated phrase, a single-word term that is the
15 combined form of the single-word terms of the hyphenated phrase, and
16 sequential subsets of the single-word terms of the hyphenated phrase;
17 h) a proper name, having more than two single-word terms, expanded
18 to include, as separate search terms, the single-word terms of said proper
19 name; and
20 i) a search term delineated by a predetermined phrase word-length.

1 8. The method of Claim 7 wherein said fielded data of each member of
2 said set of records includes descriptive information qualifying the contents of a
3 respective member of said set of databases.

1 9. The method of Claim 1, 4 or 8 wherein said document frequency is
2 calculated in relation to the number of documents in a predetermined member of said
3 set of databases that contains a predetermined one of said predetermined terms and the
4 number of documents within said predetermined member of said set of databases,
5 constrained to being in excess of a predetermined threshold number.

1 10. An method of executing a query selectively against a collection of
2 databases, each including a plurality of documents, said method comprising the steps
3 of:

4 a) receiving a first predetermined query including a set of predetermined
5 search terms from a user;

6 b) utilizing said first predetermined query to select a set of databases
7 from said collection of databases;

8 c) optionally reporting said set of databases to said user and permitting
9 said first predetermined query to be modified and substituted for said first

10 predetermined query, said steps of utilizing and optionally reporting then being
11 repeated;

12 d) optionally receiving a second predetermined query from said user and
13 substituting said second predetermined query for said first predetermined query; and

14 e) searching said set of databases to select a set of documents that are
15 responsive to said first predetermined query.

1 11. The method of Claim 10 wherein said step of utilizing includes the steps
2 of:

3 a) determining an inverse collection frequency for each member of said
4 set of predetermined search terms with respect to each said database and said set of
5 databases, and determining a document frequency for each member of said set of
6 predetermined search terms with respect to each said database;

7 b) determining a ranking value for each said database based on a sum
8 of the products of said inverse collection frequencies for said set of predetermined
9 search terms and said document frequencies for respective members of said set of
10 search terms; and

11 c) selecting said set of databases from said collection of databases based
12 on predetermined criteria dependant on said ranking value for each said database.

1 12. The method of Claim 11 wherein said of step of utilizing uses a meta-
2 index data file that stores a plurality of records that respectively represent the members
3 of a superset of said collection of databases, each of said records storing any of search
4 term data, statistical data, and fielded data descriptive of a corresponding member of
5 said superset of databases, said inverse collection frequencies and said document
6 frequencies being determinable from said search term data and said statistical data.

1 13. The method of Claim 12 further comprising the step of prequalifying

2 said collection of databases by selection from said superset of databases based on a set
3 of conditional values established by said user in connection with said set of
4 predetermined search terms, the selection of said collection of databases being based
5 on an intersection of said set of conditional values and said fielded data stored by said
6 plurality of records.

1 14. The method of Claim 13 wherein said step of obtaining includes said
2 step of prequalifying said collection of databases, and wherein said step of optionally
3 reporting includes permitting said set of conditional values to be modified and
4 substituted as part of said first predetermined query, said steps of prequalifying,
5 utilizing and optionally reporting then being repeated;

1 15. A method of selecting a subset of a set of document collections
2 containing documents to search based upon a predetermined query text including a
3 search term, said method comprising the steps of:

4 a) accessing a meta-file representative of said set of document
5 collections, including a search term occurrence list;

6 b) determining a document frequency term for said search term relative
7 to each of said document collections within said set of document collections and an
8 inverse collection frequency term for said set of document collections, said inverse
9 collection frequency term being proportional to a ratio of the number of documents in
10 said set of document collections and the number of documents in set of document
11 collections that include said search term;

12 c) determining a term ranking for each of said document collections
13 that is proportional to the respective said document frequency terms and said inverse
14 collection frequency term;

15 d) selecting said subset of said set of document collections based on the
16 relative term ranking of each of said document collections.

1 16. The method of Claim 15 wherein said meta-file includes a list of search
2 terms that occur in said set of document collections correlated to the collection and
3 document of each search term's occurrence.

1 17. The method of Claim 16 wherein said predetermined query text includes
2 a plurality of search terms, wherein said steps (a) through (c) are performed for each
3 of said plurality of search terms, and where a final ranking for each of said document
4 collections is determined as the sum of said term rankings for each of said plurality of
5 search terms, said step (d) being performed based on the relative final ranking of each
6 of said document collections.

1 18. The method of Claim 17 wherein said document frequency terms and
2 said inverse collection frequency term determined for each of said document
3 collections is scaled to increase the numerical distribution of said document frequency
4 terms while maintaining proportionality with said inverse collection frequency term.

5 19. The method of Claim 18 wherein a log function is used to scale said
6 document frequency terms and said inverse collection frequency term determined for
7 each of said document collections.

1 20. The method of Claim 16 or 18 wherein said meta-file includes a
2 plurality of records, a predetermined record being representative of a respective
3 document collection, said predetermined record including said search term occurrence
4 list and a plurality of fielded terms provided to store characteristics of said respective
5 document collection.

1 21. The method of Claim 20 wherein said plurality of fielded terms may be
2 categorized as storing statistical data and pre-search collection qualifying data.

1 22. The method of Claim 21 wherein said statistical data may include the
2 number of documents represented by said predetermined record and the number of
3 occurrences of a predetermined search term that appears in said search term
4 occurrence list, and wherein said pre-search collection qualifying data may include a
5 language type identifier, a publication type identifier, a publication name, a publication
6 date, and a cost identifier.

1 23. The method of Claim 22 wherein said predetermined query text includes
2 a post-search qualifying term, and wherein said step of selecting said subset of
3 document collections includes a step of qualifying said subset of document collections
4 selected based on said post-search qualifying term.

1 24. A method of selecting a subset of a set of document collections to
2 search based upon a predetermined query text optionally including any of a search
3 term, a pre-search qualifier, and a post-search qualifier, said method comprising the
4 steps of:

5 a) accessing a meta-information data file that includes a plurality of
6 records representing said set of document collections, each said document collection
7 representing a plurality of documents;

8 b) pre-qualifying a set of said plurality of records based upon said pre-
9 search qualifier, if any;

10 c) determining a search term frequency value for each of said pre-
11 qualified set of said plurality of records with respect to said search term if any, said
12 search term frequency values being normalized against a common factor representative
13 of the frequency of qualifying occurrences of said search term within said documents
14 of said pre-qualified set;

15 d) determining a search term ranking for each of said pre-qualified set

16 of said plurality of records based upon said frequency values and said common factor;
17 and

18 e) selecting said subset of said set of document collections to search
19 based on said search term rankings and said post-search qualifier, if any.

20 25. The method of Claim 24 wherein said document collections have a
21 searchable collection index, and wherein said method further comprises the step of
22 applying said search term against said searchable collection index to select one or more
23 predetermined documents.

1 26. The method of Claim 25 further including a step of optionally modifying
2 any of said search term, said pre-search qualifier and said post-search qualifier and re-
3 performing said steps (b) through (e).

1 27. The method of Claim 26 wherein said normalized document usage
2 frequencies for said predetermined document collection are determined in relation to
3 the frequency of occurrence of said search significant terms within said predetermined
4 document collection.

5 28. A method of performing automated collection selection relative to a
6 plurality of document collections, each including one or more documents, using a list
7 of qualified terms developed from an input query text, said method comprising the
8 steps of:

9 a) parsing said input query text to select single-word terms and
10 multiple-word phrase terms from said query text by exclusion of predetermined
11 context-free single-word terms and punctuation;

12 b) applying each such selected term against a meta-index descriptive of
13 said document collections;

14 c) determining cumulative rankings for said document collections

15 relative to each such selected term normalized against said plurality of document
16 collections; and
17 d) selecting a set of said document collections having the highest
18 relative cumulative rankings.

1 29. The method of Claim 28 wherein said step of parsing is performed to
2 select single-word terms and multiple-word phrase terms from said input query text by
3 exclusion of predetermined context-free single word terms and punctuation where said
4 selected single-word terms and multiple-word phrase terms are identified as any of:

- 5 a) a noun phrase;
- 6 b) a word tuple;
- 7 c) a capitalized word phrase; and
- 8 d) a proper name;

9 where each of said single-word terms and multiple-word phrase terms is
10 delimited by any of a capitalized word, a space, a period, and a predetermined phrase
11 word-length.

1 30. The method of Claim 29 wherein said selected single-word terms and
2 multiple-word phrase terms are further identified as any of:

- 3 a) a hyphenated phrase expanded to include, as separate terms, the
4 single-word terms of the hyphenated phrase, a single-word term that is the combined
5 form of the single-word terms of the hyphenated phrase, and sequential pairs of the
6 single-word terms of the hyphenated phrase; and

- 7 b) a proper name, having more than two single-word terms, expanded
8 to include, as a separate terms, the single-word terms of said proper name;

9 where each of said single-word terms and multiple-word phrase terms is
10 delimited by any of a capitalized word, a space, a period, and a predetermined phrase
11 word-length.

1 31. The method of Claim 28 wherein said meta-index is provided with
2 statistical data and fielded data descriptive of said document collections, said statistical
3 data including frequency-of-occurrence statistics on multiple-word phrase terms to
4 provide single-word term proximity information, and said fielded data including
5 information respectively characterizing each of said document collections.

1 32. The method of claim 31 wherein said frequency-of-occurrence statistical
2 data describes the number of occurrences of unstemmed single-word and multiple-
3 word phrase terms within a corresponding one of said document collections, and
4 wherein said step of applying includes execution of a stemming algorithm.

1 33. The method of claim 28 wherein each of said document collections are
2 represented as records in said meta-index, wherein a predetermined record stores
3 fielded data descriptive of said document collection represented by said predetermined
4 record, and a predetermined term with corresponding statistical data, said
5 predetermined term being qualified for entry into said predetermined record where:

6 a) the number of occurrences of said predetermined term within a
7 predetermined document collection is in excess of a first predetermined number;

8 b) the number of occurrences of said predetermined term within at least
9 one document within said predetermined document collection is in excess of a second
10 predetermined number; or

11 c) said predetermined term occurs within a number of documents within
12 said predetermined document collection in excess of a third predetermined number.

1 34. The method of Claim 33 wherein at least one of said first, second, and
2 third predetermined numbers is in excess of one.

1 35. The method of Claim 34 wherein said statistical data includes, for said
2 predetermined term, the number of occurrences of said predetermined term within the
3 documents of said predetermined document collection, the number of documents
4 within said predetermined document collection containing said predetermined term, the
5 number of qualified terms that occur in said predetermined document collection and
6 the number of documents within said predetermined document collection.

7 36. The method of Claim 28 wherein said step of determining cumulative
8 term rankings for said document collections includes, for each of said document
9 collections and for each of said terms relative to a respective one of said document
10 collections, the steps of:

11 a) calculating an initial term ranking for a predetermined term relative
12 to a predetermined document collection based on a ratio of the number of documents
13 having a qualified number of occurrences of said predetermined term in said
14 predetermined document collection and a qualified number of documents within said
15 predetermined document collection;

16 b) scaling said initial term ranking;

17 c) calculating a normalizing factor based on the ratio of the total
18 number of documents in said document collections and the total number of documents
19 in said document collections having a qualified number of occurrences of said
20 predetermined term;

21 d) scaling said normalizing factor;

22 e) calculating a product of said scaled initial term ranking and said
23 scaled normalizing factor to provide a term ranking for said predetermined term; and

24 f) summing said products corresponding to each of said terms relative
25 to said predetermined document collection to provide said cumulative term ranking for
26 said predetermined document collection.

1 37. The method of Claim 36 wherein said meta-index stores a plurality of

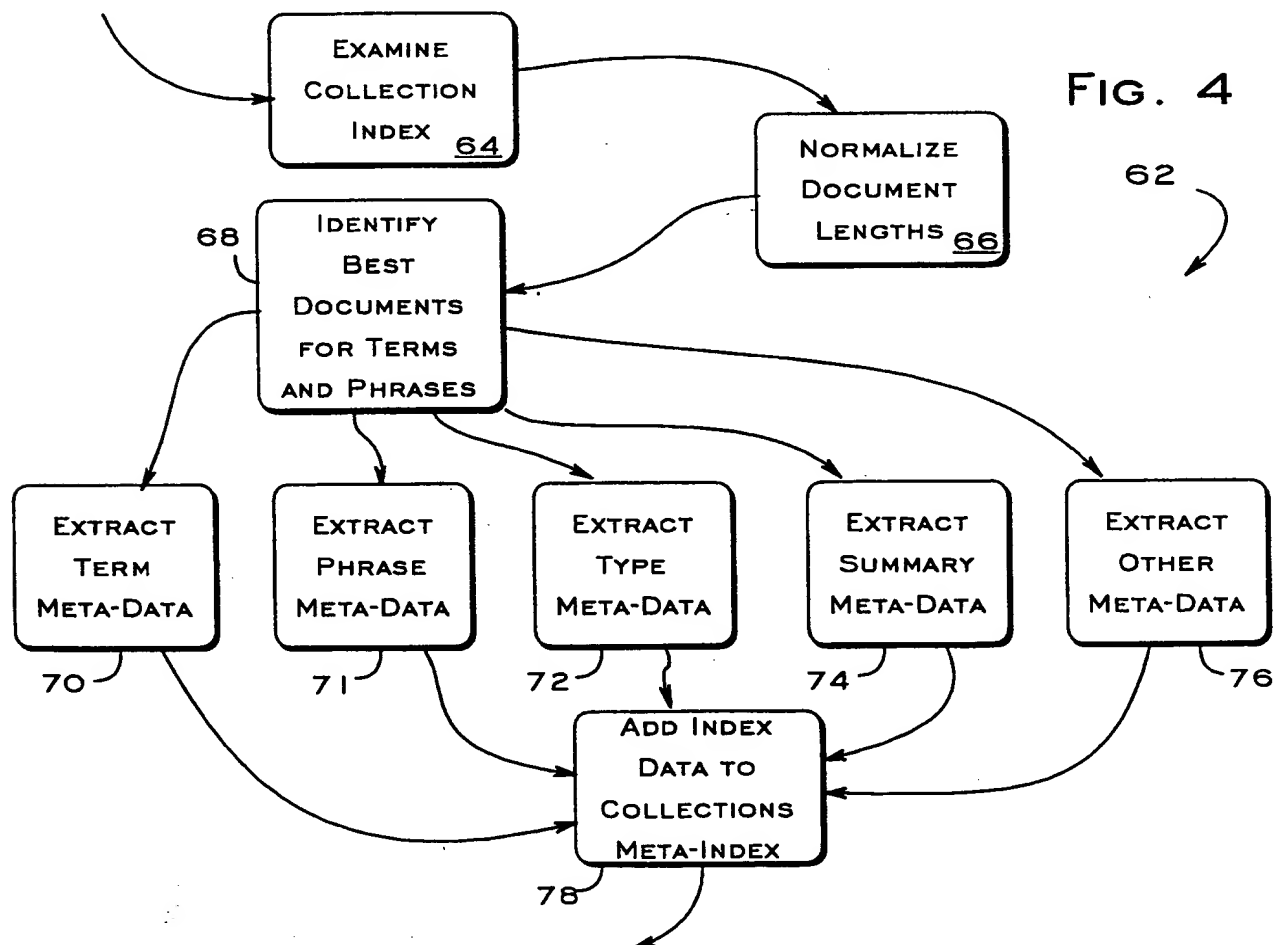
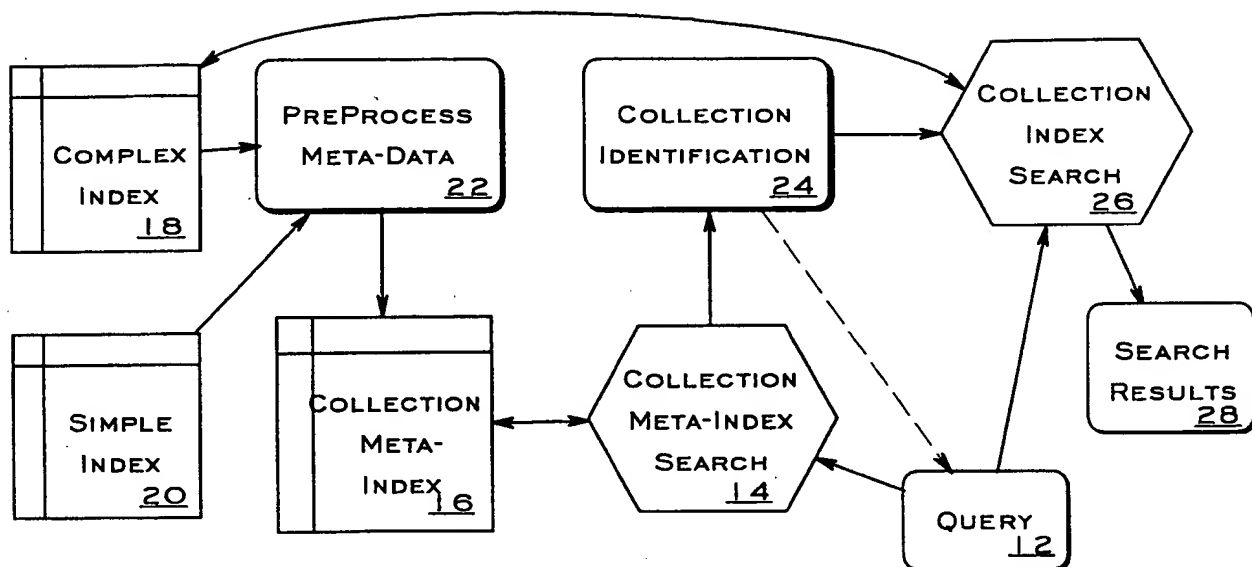
2 records representing respective said document collections, each said record including
3 a term list including term frequency of occurrence data and fielded data descriptive of
4 said respective document collection.

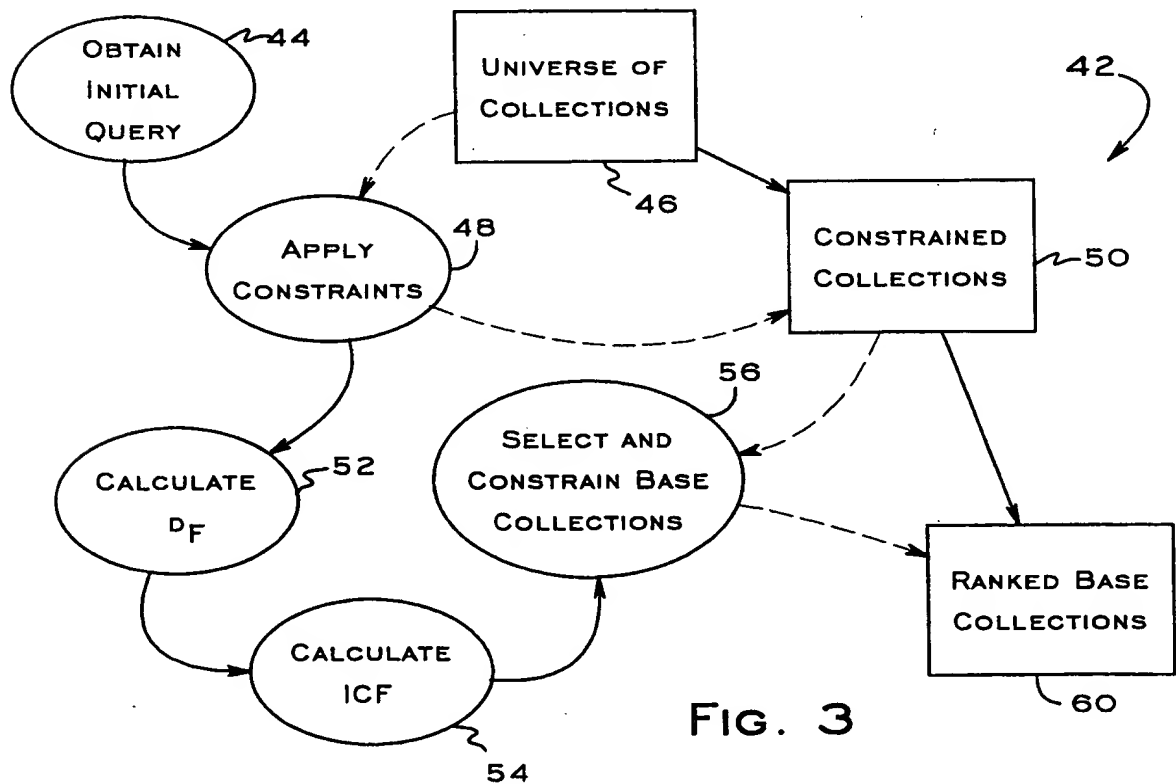
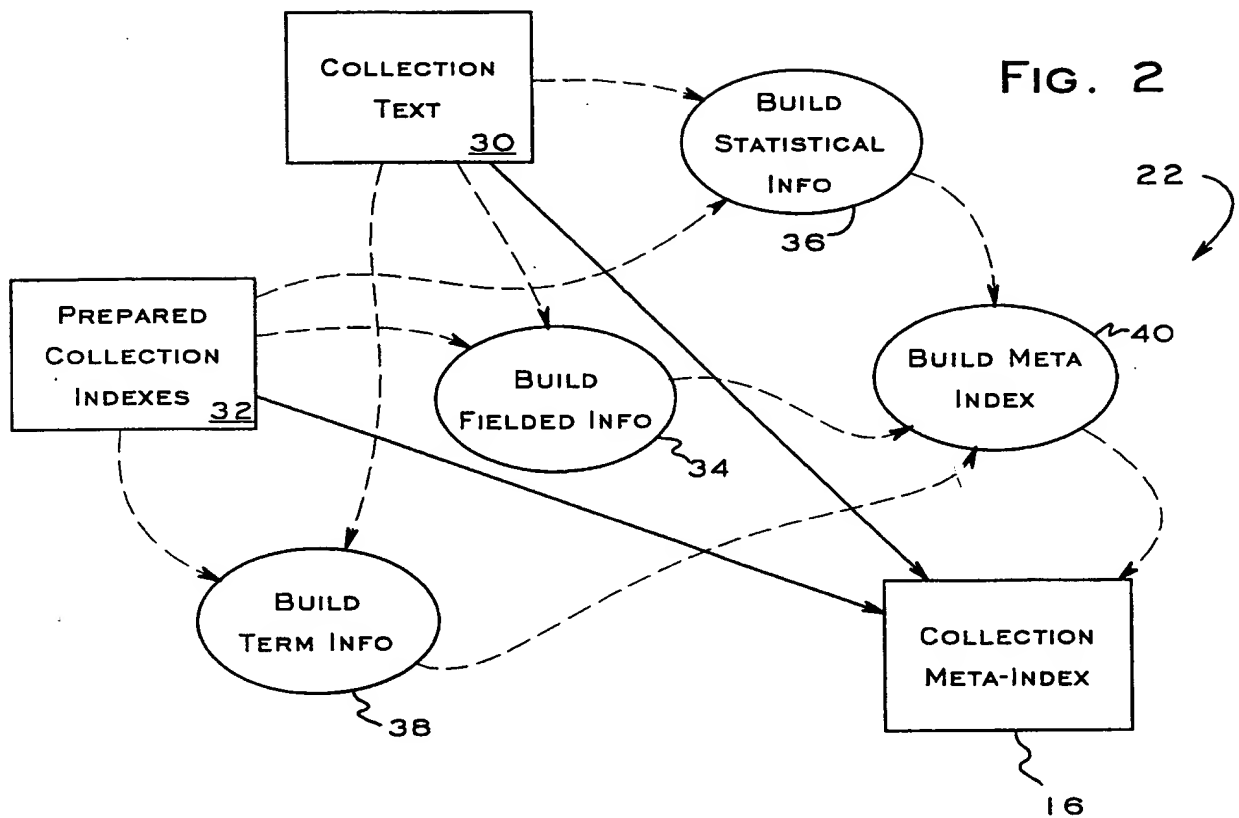
1 38. The method of Claim 37 wherein said step of parsing includes the
2 identification of pre-search qualifiers from said query text, said method further
3 including, before the step of applying, a step of pre-qualifying said document
4 collections to select a set of document collections use that meet the conditions of said
5 pre-search qualifiers, said set of document collections being said document collections
6 for subsequent steps.

1 39. The method of Claim 38 wherein said step of parsing includes the
2 identification of post-search qualifiers from said query text, said step of selecting
3 including qualifying said document collections as selected to include only those that
4 meet the conditions of said post-search qualifiers.

1 40. The method of Claim 39 wherein said step of scaling said normalizing
2 factor provides for storing said scaled normalizing factor and wherein said step of
3 calculating a product provides for storing said term rankings, said stored normalizing
4 factors and said stored term rankings being useable subsequently in said method, said
5 method further including an optional step of permitting modification of said query text
6 and resubmitting said modified query text as said query text in said step of parsing, said
7 stored normalizing factors and said stored term rankings being used as appropriate
8 where common terms and document collections are present in said query text and said
9 query text as modified to reduce time delays due to redundant calculations.

1/3





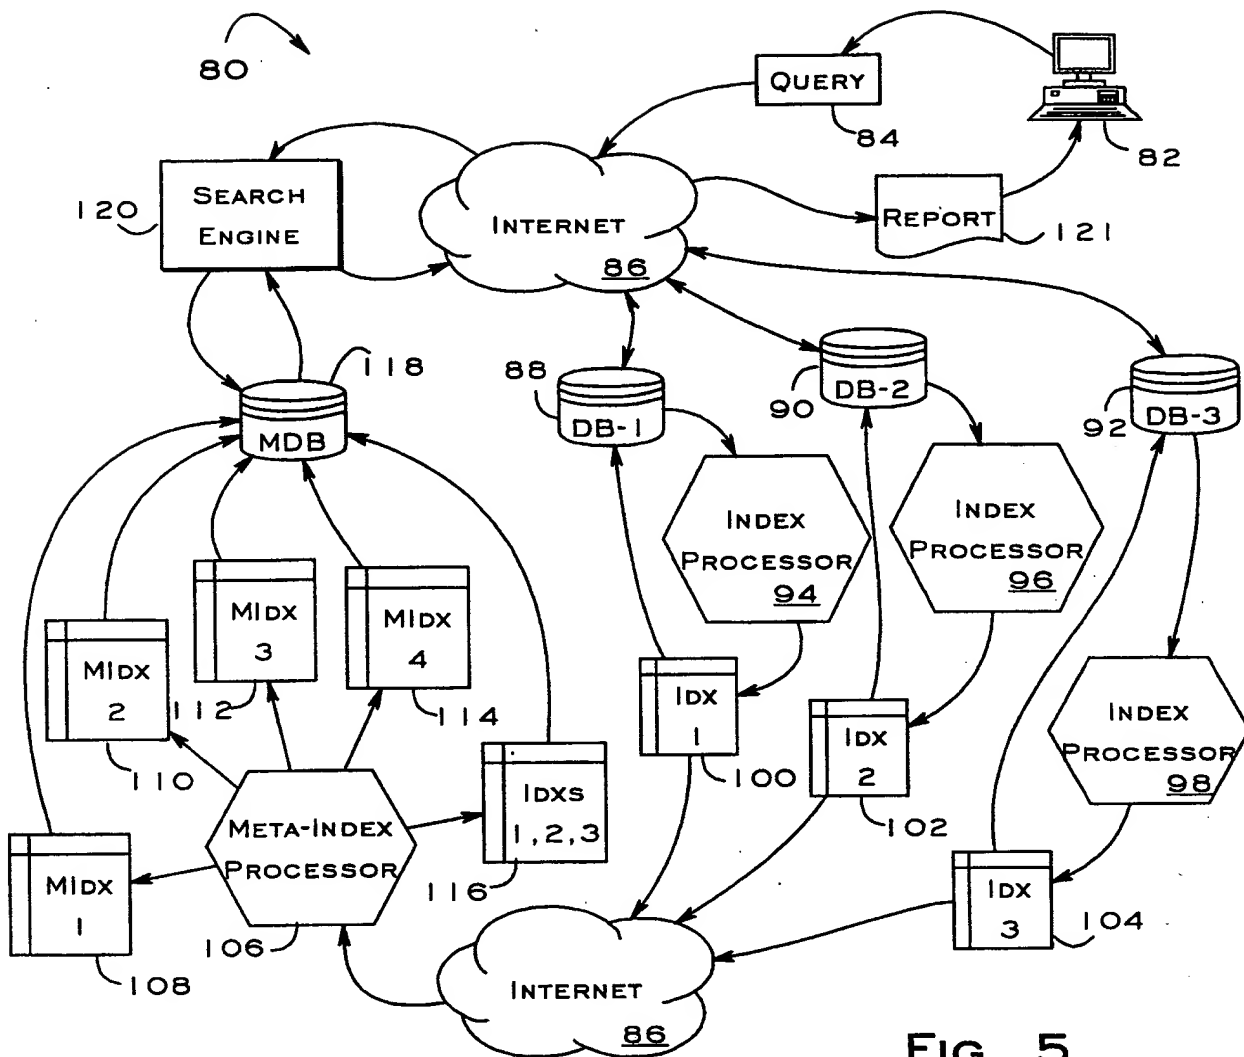


FIG. 5




FIG. 6

INTERNATIONAL SEARCH REPORT

International Application No

PCT/US 98/18844

A. CLASSIFICATION OF SUBJECT MATTER

IPC 6 G06F17/30

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC 6 G06F

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	<p>CALLAN J P ET AL: "Searching distributed collections with inference networks" 18TH INTERNATIONAL ACM SIGIR CONFERENCE ON RESEARCH AND DEVELOPMENT IN INFORMATION RETRIEVAL, SEATTLE, WA, USA, 9-13 JULY 1995, vol. spec. issue., pages 21-28, XP002085757 ISSN 0163-5840, SIGIR Forum, 1995, USA see page 21, right-hand column, paragraph 2 - page 22, right-hand column, line 35 see page 25, right-hand column, paragraph 5 - page 26, right-hand column, paragraph 6.3</p> <p style="text-align: center;">--- -/-</p>	1-40

☒ Further documents are listed in the continuation of box C.☐ Patent family members are listed in annex.

* Special categories of cited documents:

- "A" document defining the general state of the art which is not considered to be of particular relevance
- "E" earlier document but published on or after the international filing date
- "L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- "O" document referring to an oral disclosure, use, exhibition or other means
- "P" document published prior to the international filing date but later than the priority date claimed

- "T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
- "X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
- "Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.
- "S" document member of the same patent family

Date of the actual completion of the international search

25 November 1998

Date of mailing of the international search report

08/12/1998

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,
Fax: (+31-70) 340-3016

Authorized officer

Fournier, C

INTERNATIONAL SEARCH REPORT

International Application No

PCT/US 98/18844

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category	Citation of document, with indication where appropriate, of the relevant passages	Relevant to claim No.
A	<p>GRAVANO L ET AL: "Precision and recall of GLOSS estimators for database discovery" PROCEEDINGS OF THE THIRD INTERNATIONAL CONFERENCE ON PARALLEL AND DISTRIBUTED INFORMATION SYSTEMS (CAT. NO.94TH0668-4), PROCEEDINGS OF 3RD INTERNATIONAL CONFERENCE ON PARALLEL AND DISTRIBUTED INFORMATION SYSTEMS, AUSTIN, TX, USA, 28-30 SEPT. 1994, pages 103-106, XP002085758</p> <p>ISBN 0-8186-6400-2, 1994, Los Alamitos, CA, USA, IEEE Comput. Soc. Press, USA</p> <p>see page 103, left-hand column, paragraph 1 - page 104, left-hand column, paragraph 3</p> <p style="text-align: center;">---</p>	<p>1,10,15, 24,28</p>
A	<p style="text-align: center;">---</p> <p>"PARCI ARCHITECTURE"</p> <p>IBM TECHNICAL DISCLOSURE BULLETIN, vol. 38, no. 3, 1 March 1995, pages 13-18, XP000507963</p> <p>see the whole document</p> <p style="text-align: center;">-----</p>	<p>1,10,15, 24,28</p>